




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

<p>(51) International Patent Classification ⁶ : H04N 7/30</p>	<p>A1</p>	<p>(11) International Publication Number: WO 97/16028</p>	<p>(43) International Publication Date: 1 May 1997 (01.05.97)</p>
<p>(21) International Application Number: PCT/US96/17045</p> <p>(22) International Filing Date: 23 October 1996 (23.10.96)</p> <p>(30) Priority Data: 60/007,013 25 October 1995 (25.10.95) US</p> <p>(71) Applicants: DAVID SARNOFF RESEARCH CENTER, INC. [US/US]; 201 Washington Road, CN-5300, Princeton, NJ 08543-5300 (US). SHARP CORPORATION [JP/JP]; 2613-1, Ichinamoto-cho, Tenri, Nara 632 (JP).</p> <p>(72) Inventors: MARTUCCI, Stephen, Anthony; 26-12 Fox Run Drive, Plainsboro, NJ 08536 (US). SODAGAR, Iraj; 2810 Hunter Glen Drive, Plainsboro, NJ 08536 (US). ZHANG, Ya-Qin; 98 Ashford Drive, Plainsboro, NJ 08536 (US).</p> <p>(74) Agent: BURKE, William, J.; David Sarnoff Research Center, Inc., 201 Washington Road, CN-5300, Princeton, NJ 08543-5300 (US).</p>		<p>(81) Designated States: JP, KR, European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).</p> <p>Published</p> <p><i>With international search report.</i></p> <p><i>Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i></p>	

(54) Title: OVERLAPPING BLOCK ZEROTREE WAVELET IMAGE CODER

(57) Abstract

An apparatus (100) and a concomitant method for encoding video frame sequences or input images using overlapping block motion compensation (104) in conjunction with zerotree wavelet coding (106). The method partitions each input image into a plurality of overlapping blocks and applies a wavelet transform (122) on the partitioned input image.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AM	Armenia	GB	United Kingdom	MW	Malawi
AT	Austria	GE	Georgia	MX	Mexico
AU	Australia	GN	Guinea	NE	Niger
BB	Barbados	GR	Greece	NL	Netherlands
BE	Belgium	HU	Hungary	NO	Norway
BF	Burkina Faso	IE	Ireland	NZ	New Zealand
BG	Bulgaria	IT	Italy	PL	Poland
BJ	Benin	JP	Japan	PT	Portugal
BR	Brazil	KE	Kenya	RO	Romania
BY	Belarus	KG	Kyrgyzstan	RU	Russian Federation
CA	Canada	KP	Democratic People's Republic of Korea	SD	Sudan
CF	Central African Republic	KR	Republic of Korea	SE	Sweden
CG	Congo	KZ	Kazakhstan	SG	Singapore
CH	Switzerland	LI	Liechtenstein	SI	Slovenia
CI	Côte d'Ivoire	LK	Sri Lanka	SK	Slovakia
CM	Cameroon	LR	Liberia	SN	Senegal
CN	China	LT	Lithuania	SZ	Swaziland
CS	Czechoslovakia	LU	Luxembourg	TD	Chad
CZ	Czech Republic	LV	Latvia	TG	Togo
DE	Germany	MC	Monaco	TJ	Tajikistan
DK	Denmark	MD	Republic of Moldova	TT	Trinidad and Tobago
EE	Estonia	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	UG	Uganda
FI	Finland	MN	Mongolia	US	United States of America
FR	France	MR	Mauritania	UZ	Uzbekistan
GA	Gabon			VN	Viet Nam

OVERLAPPING BLOCK ZEROTREE WAVELET IMAGE CODER

This application claims the benefit of U.S. Provisional Application
5 No. 60/007,013 filed October 25, 1995.

The invention relates to video signal encoding systems and, more particularly, to apparatus and concomitant methods of encoding video signals using overlapping block motion estimation compensation and zerotree wavelet coding.

10

BACKGROUND OF THE DISCLOSURE

Standardized, block-based video signal encoding techniques such as the standards promulgated by the Moving Pictures Experts Group (MPEG) use block motion estimation compensation to remove
15 inter-frame redundancy from successive image frames within a video signal. The standards specify the use of non-overlapping rectangular blocks of image pixels from successive frames to generate a set of motion vectors that represent inter frame motion. Systems implementing these standards require coding of the motion vectors and any residual
20 differences between successive frames that are not represented by the motion vector compensation.

All operations in block-based motion compensation systems are accomplished on a block-by-block basis. As these operations are applied to each block independently, the neighboring motion vectors are not
25 always the same. Consequently, the system produces discontinuities (in the form of blocking edges) in residuals. Such discontinuities are difficult to code and can use a substantial number of coding bits.

One improvement to the standard block motion compensation systems used to address the blocking effect involves using rectangular
30 blocks of pixels that slightly overlap with one another, i.e., any two adjacent blocks overlap by one or more pixels. Such block overlap reduces the discontinuities in the residuals and improves the blocking effect.

Recently, there has been developed a coding technique that finds
35 use in coding still (photographic) images. This technique is useful at

-2-

removing intra frame information redundancy and for efficiently coding a single frame of an image. The technique is disclosed in U.S. patent 5,412,741, issued May 2, 1995 and herein incorporated by reference, which describes using the discrete wavelet transform (DWT) in deriving
5 embedded zerotree wavelets (EZW) to form an hierarchical image representation. This hierarchical representation is then efficiently coded using entropy coding.

Heretofore, the zerotree wavelet coding technique has not been combined with a overlapping block motion compensation technique such
10 that sequences of images can be efficiently coded by compressing both the inter-frame and intra-frame redundancies.

Therefore, there is a need in the art for a low bit rate video encoder using both overlapping block motion compensation and zerotree wavelet coding.

15

SUMMARY OF THE INVENTION

The present invention is an apparatus and a concomitant method of encoding video frame sequences (input images) using overlapping block motion compensation in conjunction with zerotree wavelet coding.
20 The method partitions each input image into a plurality of overlapping blocks and applies wavelet transform on the partitioned input image.

BRIEF DESCRIPTION OF THE DRAWINGS

The teachings of the present invention can be readily understood
25 by considering the following detailed description in conjunction with the accompanying drawings, in which:

FIG. 1 depicts a block diagram of an encoder of the present invention;

FIG. 2 illustrates motion vector utilization in a BMC technique;

30 FIG. 3 illustrates motion vector utilization in a OBMC technique;

FIG. 4 depicts detailed block diagram of an overlapping block motion compensator used in the present invention;

FIG. 5 illustrates quadrilateral OBMC blocks;

FIG. 6 depicts an input image that has been partitioned into polygonal shaped blocks;

FIG. 7 illustrates motion vector utilization in a OBMC technique that uses the polygonal shaped blocks of FIG. 6;

5 FIG. 8 illustrates a four stage tree structure for implementing the DWT of the present invention;

FIG. 9 illustrates a four stage tree structure for implementing the inverse DWT of the present invention;

10 FIG. 10 depicts a block diagram of one filter of the DWT shown in FIG. 8;

FIG. 11 depicts a block diagram of one filter of the inverse DWT shown in FIG. 9;

FIG. 12 depicts a mid-riser, uniform quantizer with a dead zone around zero;

15 FIG. 13 depicts a time-frequency tiling in a full wavelet packet transform;

FIG. 14 depicts a time-frequency tiling in a wavelet transform;

FIG. 15 depicts a time-frequency tiling in a wavelet packet transform; and

20 FIG. 16 depicts a time-frequency tiling in an adaptive wavelet packet transform.

To facilitate understanding, identical reference numerals have been used, where possible, to designate identical elements that are common to the figures.

25

DETAILED DESCRIPTION

FIG. 1 depicts an encoder 100 that incorporates the teachings of the present invention. The encoder contains an overlapping block motion compensator (OBMC) 104, subtractor 102, wavelet
30 generator/coder 106, bit rate controller 110, inverse wavelet generator 112, summer 114, frame buffer 116, output buffer 118 and entropy coder 120.

In general, the input is a video image (a two-dimensional array of pixels (pels) defining a frame in a video sequence). To accurately

transmit the image through a low bit rate channel, the spatial and temporal redundancy in the video frame sequence must be substantially reduced. This is generally accomplished by coding and transmitting only the differences between successive frames. The encoder has three
5 functions: first, it produces a plurality of motion vectors that represent motion that occurs between frames; second, it predicts the present frame using a reconstructed version of the previous frame and the motion vectors; and third, the predicted frame is subtracted from the present frame to produce a frame of residuals that are coded and transmitted
10 along with the motion vectors to a receiver. Within the receiver, a decoder reconstructs each video frame using the coded residuals and motion vectors.

Specifically, the present input image is sent to the bit rate controller 110, the subtractor 102 and the OBMC 104. The OBMC
15 partitions the image into a plurality of overlapping macroblocks, i.e., preferably, as shall be discussed below, into polygonal-shaped blocks of pixels. The macroblocks in the present image are compared to a previously transmitted input image. Typically, the previously transmitted input image is a previous frame in the video sequence, but it
20 may be a frame from some other portion of the sequence. The OBMC generates motion vectors representing motion that has occurred between the previous frame and the present frame. The motion vectors are coded by the entropy coder 120 such as a Huffman coder and sent to buffer 118 for ultimate transmission to a receiver. These motion vectors are then
25 applied to the previously processed image to produce a predicted image. Although the OBMC is illustrated as a single module, those skilled in the art will realize that the functions performed by the OBMC can be implemented using separate modules, e.g., a motion estimation module and a motion compensation module.

30 More specifically, motion compensation algorithms are employed to reduce the temporal redundancy of video sequences. In a conventional motion compensation scheme, the input signal is segmented into non-overlapped blocks. Using a block matching criteria such as sum absolute difference, one block from the previous frame is

chosen as the best predication for each block of the current frame. A motion vector is also obtained to represent the prediction operation. FIG. 2 shows the conventional block matching motion compensation (BMC). In this figure, the best match for the block (i,j) of frame n is found to be
 5 block (p,q) of frame $n-1$:

$$\hat{B}(n,i,j) = B(n-1,i-k,j-l) \quad (1)$$

The distance between these blocks is shown by vector $v=(k,l)$. Therefore, with the given previous frame, the current frame of the input signal can
 10 be represented by a motion vector field and a motion compensated residual (MCR) frame. The MCR frame is obtained by differencing the current frame and the predicted frame:

$$R(n,i,j) = B(n,i,j) - \hat{B}(n,i,k) \quad (2)$$

15

For compression efficiency, the motion vector field and MCR frame can be compressed instead of the entire current input frame. All operations in block matching motion compensation are accomplished on a block-by-block basis. As these operations are applied to each block
 20 independently, the neighboring motion vectors are not always the same. Consequently, a BMC-type system produces discontinuities (in the form of blocking edges) in MCR frames. Such discontinuities are difficult to code and can use a substantial number of coding bits.

In conventional block based coding schemes, a block-based
 25 transform such as DCT is applied to the MCR frames. Generally, the same block segmentation size and shape used for motion compensation, is also used for the transform computation. Therefore, the discontinuities of the MCR frames do not reduce the coding efficiency. However, these blocking edges increase the blocking effects in the
 30 decoded sequence. If, instead of DCT, a global transform such as a wavelet transform is applied to the MCR frames, the blocking edges generated by the BMC are transformed. In the transformed domain, these edges are represented with large magnitude coefficients in high frequencies and therefore reduce the coding efficiency. The transformed

discontinuities also produce the blocking effects in the decoded sequence. Therefore, an effective reduction of the blocking edges is required to improve the quality of the coding when using a wavelet transform.

One elegant solution for reducing the discontinuities, is to overlap
 5 the blocks in the motion compensation process, i.e., an overlapped block motion compensation (OBMC) approach. FIG. 3 depicts motion compensated blocks translated by motion vectors, where the blocks overlap one another. Specifically, each compensated block overlaps its eight neighboring blocks. A window function is also employed for
 10 motion compensation of the overlapped regions. In this figure, a block $B(i,j)$ of frame n is obtained by:

$$B(n,i,j) = \sum_k \sum_l W(k,l) \times B(n-1,i-k,j-l) \quad (3)$$

15 where $W(k,l)$ is the matrix representing the window function and \times represent element by element product-sum operation. FIG. 4 depicts a detailed block diagram of the OBMC 104 containing a plurality of predictors 400 and a plurality of windows 402. Each predictor 400 is connected in series with a window 402 and all the predictor/window
 20 combinations are connected in parallel to summer 114. The output of each of the windows 402 are connected to summer 404. In operation, multiple blocks from previously reconstructed frames produced by the predictors 402 are windowed and added together in order to obtain a block for the motion compensated frame. Even if two neighboring
 25 motion vectors are not the same, the discontinuities between the corresponding motion compensated blocks are eliminated because of the overlap between adjacent blocks. It has been shown that even though overlapping the motion compensated blocks might introduce some unnecessary distortion to some blocks, it significantly improves the
 30 efficiency of inter frame coding when using the wavelet transform. The optimal window shape is a function of the spatial correlation of the input signal. It has also been shown that the trapezoid or raised-cosine shape windows are near-optimal for sequences that commonly appear in

video-phone applications such as sequences depicting the head and shoulders of a caller.

The present invention uses an overlapping block motion compensation approach that uses quadrilateral overlap as shown in
5 FIG. 5. In this approach, each block is divided into four quarters, each of which overlaps with three different quarters of the neighboring blocks. The raised-cosine window function is also chosen for overlapping the blocks:

$$10 \quad W(i, j) = w(i)w(j) \quad (4)$$

$$w(i) = \frac{1}{256} \sin^2 \frac{\pi(i + .5)}{16} \quad (5)$$

for a window of size 16×16 and $i=0,1,\dots,15$. It has been observed that the performance of the quality of the wavelet coder is significantly improved
15 when the OBMC is employed rather than BMC. As an example, in the coding of 200 frames of the "mother-daughter" sequence, at the rate of 25 Kbits/sec, the average signal-to-noise ratio (SNR) has improved by 0.5 dB by using OBMC in lieu of BMC.

Note that in BMC, each block is predicted independently of the
20 other blocks. The predication quality of each block in an OBMC scheme is affected by the neighboring motion vectors. In most overlapping schemes in the prior art, these vectors are still estimated independently. One further step for improvement of the OBMC scheme is to consider the effects of each motion vector in the predication of neighboring blocks.
25 One approach uses an iterative algorithm for estimation of motion vectors. An iterative approach has been shown that 2 or 3 iterations of motion estimation can reduce the prediction error up to 15% percent.

Preferably, the blocks defined by the OBMC are polygonal in shape, e.g., having a hexagonal shape. The hexagonal shape more
30 closely represents human visual properties than a square block. As such, the use of hexagonal blocks reduces visual blocking effect. Additionally, hexagonal blocks form a better representation of human

visual properties. FIG. 6 depicts a illustrative input frame 600 that has been partitioned into a plurality of non-overlapping hexagonal blocks 602.

FIG. 7 depicts an illustrative motion vector 700 that relates
5 block 702 in frame i to block 702 in frame $i-1$ and depicts an illustrative motion vector 704 that relates block 706 in frame i to block 706 in frame $i-1$. Each motion vector ensures that the blocks overlap one another.

Once the motion vectors are available, they are used to produce the predicted frame (on path 134 in FIG. 1). The predicted frame is, in turn,
10 used to produce an MCR frame. Subtractor 102 subtracts, pixel-by-pixel, the predicted frame from the present input frame. The result is a frame of motion compensated residuals (an MCR frame).

The wavelet generator/coder 106 processes the residuals using a conventional discrete wavelet transform algorithm and a zerotree coding
15 and quantization technique. The wavelet generator/coder contains a discrete wavelet transform (DWT) block 122, a quantizer (Q) 124, a zerotree coder 126 and an entropy coder 128. A detailed discussion of a wavelet generator/coder is disclosed in U.S. patent 5,412,741 issued May 2, 1995 and incorporated herein by reference. An alternative zerotree
20 encoder is disclosed in provisional application serial number 60/007,012, filed October 25, 1995, Attorney Docket Number 11908 (converted into US patent application number _____, filed _____, Attorney Docket Number DSRC 11908) and also incorporated herein by reference.

The inverse wavelet generator 112 contains an inverse
25 quantizer (Q^{-1}) 130 and an inverse DWT (DWT^{-1}) block 132. As such, the inverse wavelet generator attempts to reconstruct the MCR frame that forms the input to the wavelet generator/coder. This reconstructed frame of residuals is combined, in summer 114, with the predicted image to provide an estimation of the present input image. This image
30 is stored in the frame buffer 116 to be used as the previous image by the OBMC when generating the new predicted frame for the next input frame.

FIGs. 8 through 11 depict the multi-staged filter structure of the discrete wavelet transform (DWT). The DWT 122 is usually

implemented using an octave-band tree structure containing a plurality of filters 800, 802, 804, and 806. Each stage (filter stage) of the tree structure is a 2-band decomposition and only the low frequency band of this 2-band system is further decomposed. For example, FIG. 10 depicts
5 a detailed block diagram of one stage 800 (in one-dimension) of the DWT of FIG. 8. The stage 800 contains a low band filter 1002 and a high band filter 1004. Each filter is followed by a downsampler 1006 and 1008 that subsamples the filtered image.

FIGs. 9 and 11 depict the inverse DWT 132 implemented in a
10 similar structure using the corresponding two-band synthesis sections 900, 902, 904, and 906. Each stage contains two upsamplers 1106 and 1108 followed by a low band filter 1102 and a high band filter 1104. The output signals from the filters are added in summer 1116. If the basic two-band system provides perfect reconstruction, the DWT is
15 invertible. The desired DWT can be obtained by designing the two-band filter bank with the required properties.

Similar to the 1-D case, multidimensional discrete wavelet transforms are usually implemented in the form of hierarchical tree structures of filter banks. Each stage of the tree structure is a
20 multidimensional filter bank. The generalization of filter bank to multiple dimensions is straight-forward for the separable case. Separable filter banks are those which use both separable filters and separable down-upsampling functions. The implementation of separable filter banks is very efficient due to the fact that the
25 decomposition is applied in each dimension separately. Although separable filter banks do not necessarily maximize the coding gain of the transform for a given image, their performance in a wavelet structure is very near optimal for natural images. Therefore, the preferred embodiment of the invention uses separable filter banks in the
30 implementation of the 2-D wavelet transform. The subband decomposition is first applied to the rows and then the columns of pixels within an image. Therefore, using a simple recursive routine, the input image can be decomposed into four subimages in each recursion. Although the preferred embodiment of the invention uses separable

filters, an alternative embodiment uses non-separable filters or a combination of separable and non-separable filters.

One important parameter is the choice of filters in the DWT. To facilitate image compression, the need of linear filters is well known since without linear phase, the phase distortion around edges is very visible. Therefore, the invention uses linear phase filters only. Orthogonality is another issue in the filter bank. Since the decorrelation of the subband signal is a desired property in compression system, the invention uses orthogonal or near-orthogonal filter banks. Perfect reconstruction property is another important characteristic of the filter bank. So, the invention uses a near perfect reconstructing filter bank, because at low bit rates, the reconstruction error is negligible compared to the quantization errors. As for the size of filters, long filters usually tend to spread coding noise. Therefore, the use of long filters in low-bit rate coding systems results in ringing effects around sharp edges. Thus, when the invention is used in an application that uses very low bit rates, the invention generally uses short filters, although they might not have very good stop band rejection. Hence, the filters used in the invention have a length of 12-taps or less.

To demonstrate the practicality of using different filter types for different wavelet subbands, the invention was used to code the mother-daughter sequence with two different sets of wavelets: namely, two QMF9 filters and two Harr filters (See FIGs. 8-11). It was expected that the Harr wavelet would perform better on MCR frames because of its superior time-resolution compared to the QMF9 wavelet. It turned out that the overall quality of the decoded inter frame is in fact better when the Harr wavelet is used. The improvement in luminance SNR is about 0.17 dB for the mother-daughter sequence. The use of Harr filters reduces the ringing artifacts around the edges and therefore improves the subjective quality of the decoded sequence significantly. Although this transform is more blocky than the QMF9.

In the present invention, the wavelet transform is applied to the motion compensated residual (MCR) frames. Then, wavelet tree coding

is applied to the coefficients within a wavelet tree representation of each MCR frame.

Returning to FIG. 1, the discrete wavelet transform 122 is followed by a quantizer 124 that quantizes the wavelet transform coefficients. The quantizer is followed by a zerotree coder 126 and an entropy coder 128 which losslessly code the quantized coefficients. The discrete wavelet transform is a lossless and reversible process that does not yield any compression but prepares the data for the subsequent compression stages. Substantial compression, and therefore loss, occurs in the quantizer. The zerotree and entropy coders that follow the quantizer form a lossless compression stage that is used to encode the quantized coefficients with as few bits as possible.

The first embodiment of the present invention uses the Embedded Zerotree Wavelet (EZW) encoding process to perform the combined operations of wavelet transform, quantization, and zerotree/entropy coding. A detailed description of an EZW encoder can be found in U.S. patent 5,412,741 issued May 2, 1995 and incorporated herein by reference. A second (alternative) embodiment of the invention uses a different form of zerotree encoder that is disclosed in U.S. provisional patent application serial number 60/007,012, filed October 25, 1995, Attorney Docket Number 11908, entitled "APPARATUS AND METHOD FOR ENCODING ZEROTREES GENERATED BY A WAVELET-BASED CODING TECHNIQUE" (converted into US patent application number _____, filed _____, Attorney Docket Number DSRC 11908), which is herein incorporated by reference. Each form of wavelet tree encoder is discussed below.

After computing the wavelet transform, the EZW process continues by simultaneously quantizing coefficients, building zerotrees of the coefficients, and entropy coding the zerotrees and the coefficient values using an arithmetic coder. This is accomplished iteratively, where the quantization becomes finer, and therefore coding quality improves after each iteration. The iterations continue until the bit allocation for a particular frame is exhausted. The bit allocation is set by the bit rate controller 110 of the inventive encoder.

The quantization that is implemented by EZW can be characterized as a family of quantizers, each of which is a mid-riser uniform quantizer with a dead zone around zero. An illustrative example of such a quantizer is plotted in FIG. 12. After each iteration of the EZW process, all coefficients will have been effectively quantized by using one of these quantizers. As the algorithm proceeds through the next iteration, the effective quantization for each coefficient becomes that of the next finer quantizer in the family of quantizers until, at the end of the iteration, the quantization of all coefficients is that of this new quantizer. If the bit allocation is used up at the end of an iteration, all coefficients will have been quantized according to the same quantizer. It is more likely, however, that the bit allocation is exhausted before an iteration is completed, in which case the final effective quantization for each coefficient will be according to one of two quantizers, depending upon where in the scan of coefficients the algorithm stops.

Each iteration of the EZW process is characterized by a threshold, and the quantizer effectively implemented is a function of that threshold. As the iterations proceed, the threshold decreases and the quantization becomes finer. All thresholds are powers of two. The initial threshold T_d is set at that power of two such that the magnitude of at least one coefficient lies between T_d and $2*T_d$, and no coefficient has a magnitude greater than $2*T_d$. As each coefficient with magnitude c is processed by this iteration, its quantized magnitude c_q becomes:

$$c_q = \begin{cases} \text{int}[c/T_d] * T_d + T_d/2 & \text{if } c \geq T_d \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

where $\text{int}[]$ means "take the integer part of".

After this iteration, a new threshold T_s is used, where $T_s = T_d/2$. Coefficients are scanned again, and become quantized according to:

$$c_q = \begin{cases} \text{int}[c/T_s] * T_s + T_s/2 & \text{if } c \geq 2T_s \\ 0 & \text{otherwise} \end{cases} \quad (7)$$

Notice that the only difference between this quantizer and the first one is that the dead zone around zero extends to the threshold in the first quantizer but to twice the threshold in the second.

After this iteration, the threshold T_d is again used, where here T_d is set to $T_d = T_s$. Coefficients are scanned again, and become quantized according to Equation 6. After that iteration, T_s is set and Equation 7 applies. This changing of quantizers continues until the bit allocation for the frame is exhausted.

To encode the quantized coefficients, EZW uses a structure called a zerotree to order the coefficients followed by an arithmetic coder to do the actually encoding. The zerotree and arithmetic coder work together to losslessly compress the quantized coefficients to their entropy. The zerotree is regrown with each iteration of EZW so that it reflects the current state of quantization of the wavelet coefficients. There is a node in the tree for each wavelet coefficient. In U.S. patent 5,412,741, one of four symbols describes the coefficient at that node: POSITIVE SIGNIFICANT, NEGATIVE SIGNIFICANT, ZEROTREE ROOT, or ISOLATED ZERO. The zerotree coding yields compression because the ZEROTREE ROOT symbol is used to prune the tree, i.e., the branch extending from the ZEROTREE ROOT is not coded. Such pruning enables the invention to limit the number of coefficients that must be coded, thereby saving many bits.

Alternatively, in U.S. provisional patent application serial number 60/007,012 (Attorney Docket 11908) (converted into US patent application number _____, filed _____, Attorney Docket Number DSRC 11908), the zerotree encoder uses a depth first pattern for traversing the zerotree, i.e., each branch of the tree, from parent to child to grandchild and so on, is fully traversed before a next branch is traversed. The depth first tree traversal pattern is used to quantize the coefficients of the tree as well as to assign symbols to the quantized coefficients without using an iterative process such as that used in the EZW approach. Quantization can be accomplished with any quantization method that produces discrete values from a continuous wavelet coefficient value. Such quantization processes include the

-14-

mid-riser quantizer of FIG. 12. Additionally, this zerotree encoder assigns one of three symbols to each node: ZEROTREE ROOT, VALUED ZEROTREE ROOT, and VALUE. As with the EZW approach, the zerotree coding yields compression because the ZEROTREE ROOT
5 symbol is used to prune the tree, i.e., the branch extending from the ZEROTREE ROOT is not coded. Such pruning enables the invention to limit the number of coefficients that must be coded, thereby saving many bits.

The arithmetic (entropy) coder 126 that follows the zerotree coder
10 is used to actually encode the remaining symbols and the bits of the quantized coefficients. An arithmetic coder is an entropy coder that can encode a stream of symbols at their entropy. An arithmetic coder is similar to a Huffman coder, but can outperform the Huffman coder because the arithmetic coder can assign fractional bits to the symbols
15 and because the arithmetic coder can adapt more quickly to changing symbol statistics than the Huffman coder.

The particular arithmetic coder used is a direct implementation of the coder developed by Witten, Neal, and Cleary. A full description of the coder can be found in Witten et al., "Arithmetic Coding For Data
20 Compression," Communications of the ACM, Vol. 30. No. 6, pp. 520-540, June 1987. The coder operates by assigning bits to symbols according to a model of the probabilities of the symbols. An adaptive model is used which means that the probabilities are computed as the symbols are encountered in the encoding process. In this way, the data is encoded in
25 as few bits as possible and can still be decoded because the decoder can build the same model as the encoder and track the probabilities. By using an adaptive model, the arithmetic coder adapts to the changing statistics of the signal as the encoding proceeds.

The bit rate controller 110 monitors the input image, the buffer 118
30 status, the status of the entropy coder 120, and the residuals to determine an optimal quantizer scale for the quantizer 124. Control of the quantizer scale controls the ultimate number of bits that are used to code a given image frame.

Bit allocation involves decisions about how to spend the available bitrate when compressing a video sequence. Bit allocation is accomplished for two reasons in video compression. First, it is used to regulate the number of bits produced during compression. In this case, bit allocation is sometimes referred to as bit rate control. Bit rate control is important for maintaining the desired bit rate, as well as for maintaining the proper fullness in any buffers in the encoder. Second, bit allocation is used to regulate the quality of the compressed video, so as to achieve the best possible decompressed video. In this context, "best" is the decompressed video that is best suited for the purpose for which it will be used. Often, decompressed video is simply used for viewing by people, with no particular task in mind. In this case, "best" simply means "looks the best to a human observer." In other cases, however, it may be that one spatio-temporal section of the video is more important than other sections. For example, in the video-telephone application, the face is usually important, and the background is less important. In this case, "best" means the decompressed video that yields the best looking face, with much less consideration given to the quality of the background. The definition of "best" for the application of interest should be integrated into the bit allocation scheme.

There are several levels at which bit allocation can be exercised. At the highest level, bit allocation for video deals with the trade-off between spatial quality and temporal quality. This trade-off is made by varying the number of frames/second (f/s). This type of bit allocation is frame rate control. As an example, say for a given bit rate, X bits/second (b/s), one chose to use Y f/s. This would result in an average of $X/Y = Z$ bits/frame (b/f) allocated for representing each compressed frame. Instead, say that for the same bit rate X b/s, one chose to use $Y/2$ f/s. This would result in an average of $2*Z$ b/f allocated for representing each compressed frame. In the first case, a larger portion of the bitrate is being spent to represent the temporal information in the video sequence (more f/s, fewer b/f). In the second case, a larger portion of the bitrate is being spent to represent spatial information in the video sequence (fewer f/s, more b/f). Achieving the best balance, as measured by decompressed

video quality, between f/s and b/f is the goal of frame rate control. The best balance between f/s and b/f at any particular moment in time depends strongly on the spatial and the temporal complexity of the video at that moment in time. For video with relatively small amplitude
5 motion, few f/s are needed to render the motion smoothly. On the other hand, for video with relatively large amplitude motion, a relatively larger number of f/s are needed in order to represent the motion smoothly.

At the next level, bit allocation deals with the division of the
10 available bits among frames, in order to equalize spatio-temporal quality between those frames. This is "inter-frame" bit allocation. Since spatio-temporal complexity can vary widely from frame to frame, different frames might require different portions of the available bits, in order to achieve even quality among the frames. The allocation of bits for a
15 particular frame will depend strongly on the spatio-temporal complexity of that frame. Note that inter-frame bit allocation can be achieved by continuously varying the frame rate. However, even for a fixed frame rate, different allocations of bits for different frames in the video sequence may be desired.

20 At the lowest level, bit allocation involves decisions about the allocation of the available bits within a frame, in order to equalize quality spatially across the frame. This is "intra-frame" bit allocation. Since the spatial complexity of a frame can vary widely depending on spatial location, different parts of a frame might require different portions of the
25 available bits, in order to achieve even quality across the frame. As with inter-frame bit allocation, for intra-frame bit allocation, the best allocation of bits for a particular spatial location in a frame depends strongly on the spatial complexity of the frame at that location.

One bit allocation approach that can be used in the invention is a
30 basic first-pass technique. To control spatial quality, the invention uses variable spatial amplitude quantization. In this basic approach, the bit rate and frame rate are specified when the encoder is run; and they are fixed throughout the encoding of the sequence. By fixing the frame rate and the bit rate, the invention also fixes the number of bits per frame.

-17-

For each frame, the encoder selects a quantizer scale for that frame that will produce exactly the number of bits allocated for the frame (the same number of bits for every frame). In this way, the quantizer scale only varies from frame to frame.

5 In order to balance spatial and temporal quality in the decompressed video, the invention can use a frame rate control approach to augment the bit rate control approach discussed above. To facilitate the frame rate control approach, measures are needed for spatial and temporal complexity. The temporal complexity measure
10 takes into account the amplitude of the motion vectors for a frame (or set of frames), as well as higher-order statistics of the motion vector field(s). The temporal complexity measure also includes statistics associated with the motion-compensated prediction error for a frame (or set of frames). The spatial complexity measure uses statistics from the source
15 video frame(s), as well as statistics from the motion-compensated prediction error frame(s). These statistics are computed separately over different spatio-temporally localized regions in the frame(s), or over entire frame(s). Similar spatial and temporal complexity measures are also used to determine inter-frame bit allocation.

20 In order to equalize quality within a frame, the invention uses an intra-frame bit allocation approach which uses spatial amplitude quantization to regulate spatial quality. For intra-frame bit allocation, spatial segmentation of the frame is performed. This segmentation takes into account two different types of information. First, the video
25 data is classified based on relative importance for the application. Re-using a previous example, the face in a video-telephone sequence may be segmented from the background, and compressed with more fidelity (using a larger proportion of the bits allocated to the frame). The segmentation also takes into account human visual system properties,
30 and classify image regions based on the degree to which compression artifacts are masked in those regions. This segmentation information is also used to determine inter-frame bit allocation.

Although the present invention operates, as described above, using a conventional wavelet transform, other, more exotic, wavelet

transform and coding processes are available and are within the scope of the present invention. For example, wavelet packets can be used as a substitute for the convention wavelet transform. Wavelet packets differ from conventional wavelets in that bands other than the low band can be further decomposed. In a conventional wavelet transform, which is typically implemented as an octave-based filter bank, the signal is split into two bands (high and low) and then the low band is split again into a high and low band and so on. This form of filter structure is shown in FIG. 8. In a wavelet packet implementation, the high band output, instead of the low band output, can be split into a low band and a high band. Additionally, both low and high band outputs are split into low and high bands for each. This form of filter structure leads to various combinations of splitting of the bands at each stage.

This flexibility provides a wide range of possible decompositions, each of which has a different time-frequency tiling and therefore different time-frequency localization. One trivial case is the decomposition of both bands in each stage of tree structure, which is equivalent to a uniform filter bank. By using different splittings, a transform can achieve different tilings of the time-frequency plane. Four examples of complex time-frequency tiling are depicted in FIGs. 13-16. Where FIG. 13 depicts a time-frequency tiling in a full wavelet packet transform; FIG. 14 depicts a time-frequency tiling in a wavelet transform; FIG. 15 depicts a time-frequency tiling in a wavelet packet transform; and FIG. 16 depicts a time-frequency tiling in an adaptive wavelet packet transform.

In the embedded zerotree coding approach, each child can have only one parent. That means the zero-tree coding scheme cannot be directly applied to the general wavelet packet decomposition, because in a wavelet packet a child may have multiple parents. In one embodiment of the invention, the invention uses a restricted wavelet packet decomposition in which, each child can have only one parent. Note that in this embodiment, unlike the conventional wavelet transform, the parent may contain higher frequency coefficients of the signal. Using different adaptive approaches for each given MCR frame, one can find

-19-

the best (restricted) wavelet packet of that frame and then apply the embedded zero-tree coding to the wavelet packet coefficients. The 'best' transform is the one in which most significant samples are gathered in a few coarse bands. The invention can use one of two different criteria

5 for wavelet packet decomposition of the MCR. The first criterion is the energy of the subband. In this embodiment of the invention, in each stage of the tree structure the subband containing more energy than the other subbands is further decomposed. Experimental results have shown that such adaptive decomposition results in a coding quality

10 improvement. The second criterion is the pre-threshold energy. In this approach, each subband is first clipped by a constant threshold and then the energies of the subbands are compared to one another. Experimental results have shown an improvement in coding quality over the original wavelet transform.

15 Although various embodiments which incorporate the teachings of the present invention have been shown and described in detail herein, those skilled in the art can readily devise many other varied embodiments that still incorporate these teachings.

-20-

What is claimed is:

1. Apparatus for encoding an input image which is partitioned into a plurality of overlapping blocks, said apparatus comprising:
5 an overlapping block motion compensator for computing a plurality of motion vectors for said plurality of overlapping blocks and for generating a predicted image using said plurality of motion vectors;
a wavelet transform module, coupled to said overlapping block motion compensator, for applying a wavelet transformation to a
10 difference signal between the input image and said predicted image, where said wavelet transformation produces a plurality of coefficients;
a quantizer, coupled to said wavelet transform module, for quantizing said plurality of coefficients; and
a coder, coupled to said quantizer, for coding said plurality of
15 quantized coefficients.

2. The apparatus of claim 1, wherein said overlapping block motion compensator comprises:
a plurality of predictors for predicting a plurality of non-
20 overlapping blocks;
a plurality of windows, where each of said plurality of window is coupled in series to one of said plurality of predictors, for windowing said plurality of non-overlapping blocks; and
a summer, coupled to said plurality of windows, for adding said
25 plurality of windowed non-overlapping blocks to produce an overlapping block.

3. The apparatus of claim 1, wherein said wavelet transform module comprises a plurality of separable filters that are organized in a
30 tree structure having a plurality of wavelet subbands.

4. The apparatus of claim 1, further comprising:
a rate controller, coupled to said quantizer, for selecting a quantizer scale for each of said plurality of coefficients.

-21-

5. The apparatus of claim 1, further comprising a zerotree coder, coupled to said quantizer, for selectively pruning said plurality of quantized coefficients.

5 6. Method for encoding an input image which is partitioned into a plurality of overlapping blocks, said apparatus comprising:
 computing a plurality of motion vectors for said plurality of overlapping blocks;
 generating a predicted image using said plurality of motion
10 vectors;
 applying a wavelet transformation to a difference signal between the input image and said predicted image, where said wavelet transformation produces a plurality of coefficients;
 quantizing said plurality of coefficients; and
15 encoding said plurality of quantized coefficients.

7. The method of claim 6, wherein said plurality of overlapping blocks are polygonal-shaped blocks.

20 8. The method of claim 6, wherein each of said plurality of overlapping blocks is generated by applying a window function.

 9. The method of claim 6, wherein said wavelet transform applying step applies selectively different types of filtering to different
25 wavelet subbands resulting from said wavelet transform.

 10. The method of claim 6, wherein said wavelet transform applying step applies wavelet packet transform.

1/7

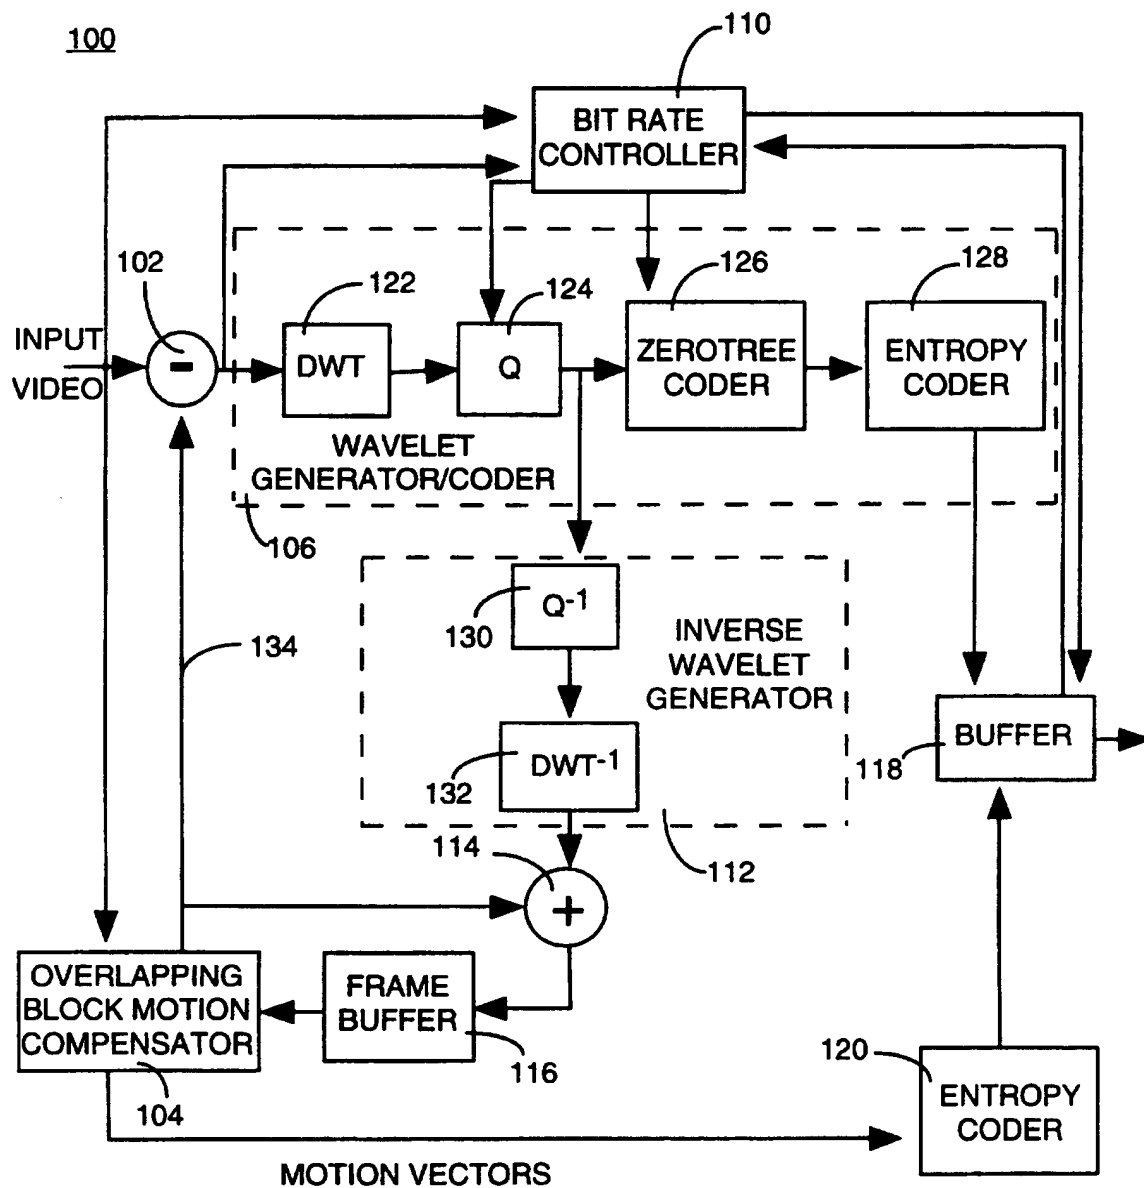


FIG. 1

2/7

FIG. 2





FIG. 3



3/7




FIG. 4




FIG. 5




FIG. 6




FIG. 7

5/7




FIG. 8




FIG. 9

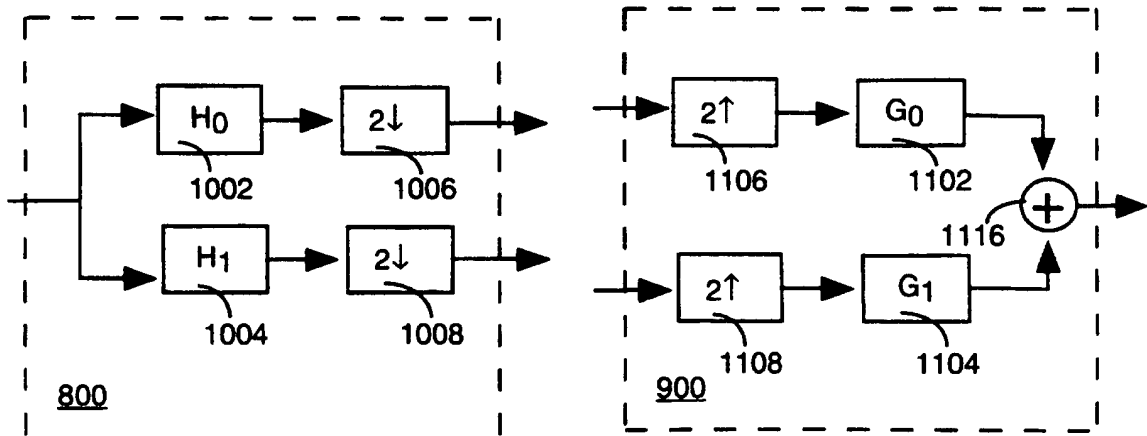


FIG. 10

FIG. 11

6/7




FIG. 12

7/7

FIG. 13




FIG. 14




FIG. 15





FIG. 16



INTERNATIONAL SEARCH REPORT

International application No.
PCT/US96/17045

A. CLASSIFICATION OF SUBJECT MATTER

IPC(6) :H04N 7/30

US CL :348/398, 402, 403, 405, 413

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 348/384, 390, 397, 398, 400-405, 409, 411-413, 415, 416; 382/232, 236, 238, 248, 251

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y,P	US 5,495,292 A (ZHANG ET AL) 27 February 1996, the whole document.	1 and 3-10
Y	US 5,311,310 A (JOZAWA ET AL) 10 May 1994, col.5, line 60 to col.13, line 12.	1 and 3-10
Y	US 5,412,741 A (SHAPIRO) 02 May 1995, col.3, line 38 to col.15, line 56.	5
A	US 5,253,058 A (GHARAVI) 12 October 1993, the whole document.	1-10
A,P	US 5,477,272 A (ZHANG ET AL) 19 December 1995, the whole document.	1-10
A	US 5,355,378 A (OHTA) 11 October 1994, the whole document.	1-10



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents:	*T	later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
A document defining the general state of the art which is not considered to be of particular relevance	*X*	document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
E earlier document published on or after the international filing date	*Y*	document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	*Z*	document member of the same patent family
O document referring to an oral disclosure, use, exhibition or other means		
P document published prior to the international filing date but later than the priority date claimed		

Date of the actual completion of the international search

21 JANUARY 1997

Date of mailing of the international search report

05 MAR 1997

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer
Bryan S. Tung
BRYAN S. TUNG

Telephone No. (703) 305-4700